

APPENDIX K - SNAPSHOTS

The GCCS Core Database uses Oracle7 snapshots to replicate the data found in 42 reference file tables. This appendix gives an overview of the internal objects involved, how they were created, and some of the related administrative tasks.

Further information on Oracle7 snapshots can be found in Oracle manuals such as *Oracle7 Server Concepts Manual*, *Oracle7 Server Administrator's Guide*, and *Oracle7 Server Distributed Systems: Replicated Data*. For more information specific to the GCCS Core Database, refer to the readme text files **SNAPSHOT.txt** and **SNAP_ADMIN.txt** located under the /h/SMDB Segment.

K.1 SNAPSHOT ARCHITECTURE

Simple, read-only snapshots are currently implemented in the GCCS Core Database to propagate changes to the reference files. They are considered to be simple snapshots because each snapshot site has a one-to-one copy of each master table. They are read-only at all the sites except the master site. Since NMCC is the master site, it is the only database through which data in these 42 tables can be modified.

K.1.1 Internal Oracle Objects

Oracle uses various internal objects to keep each snapshot in sync with the corresponding master table. All of these objects are maintained by Oracle and should never be manually modified in any way. A subset of Oracle internal objects exists at each site depending on whether the site is a snapshot site or the master site. The following objects are unique to the snapshot sites:

- Ⓒ Views named MVIEW\$_snapshot_name created by Oracle to reflect the most current data from the master tables.
- Ⓒ Base tables named SNAP\$_snapshot_name created by Oracle to store the rows retrieved by the snapshot's defining query.
- Read-only views named snapshot_name which the users query against.
- Ⓒ Indexes named I_SNAP\$_snapshot_name created by Oracle on the base tables.

The following objects are unique to the master site:

- Ⓒ Master tables for each reference file.
- Ⓒ Domain, range, and referential integrity constraints.

- C Snapshot log tables named MLOG\$_master table name created by Oracle to store the ROWID, time stamp, and data manipulation type code (ie - insert, update, delete) of rows in the master table that have been updated since the last refresh.
- C AFTER ROW triggers named TLOG\$_master table name created by Oracle on the master table to insert the ROWID and time stamp of inserted, updated, and deleted rows into the snapshot log table.

Note: That the snapshot name is always the same as the corresponding master table name. Figure K-1 shows an example.

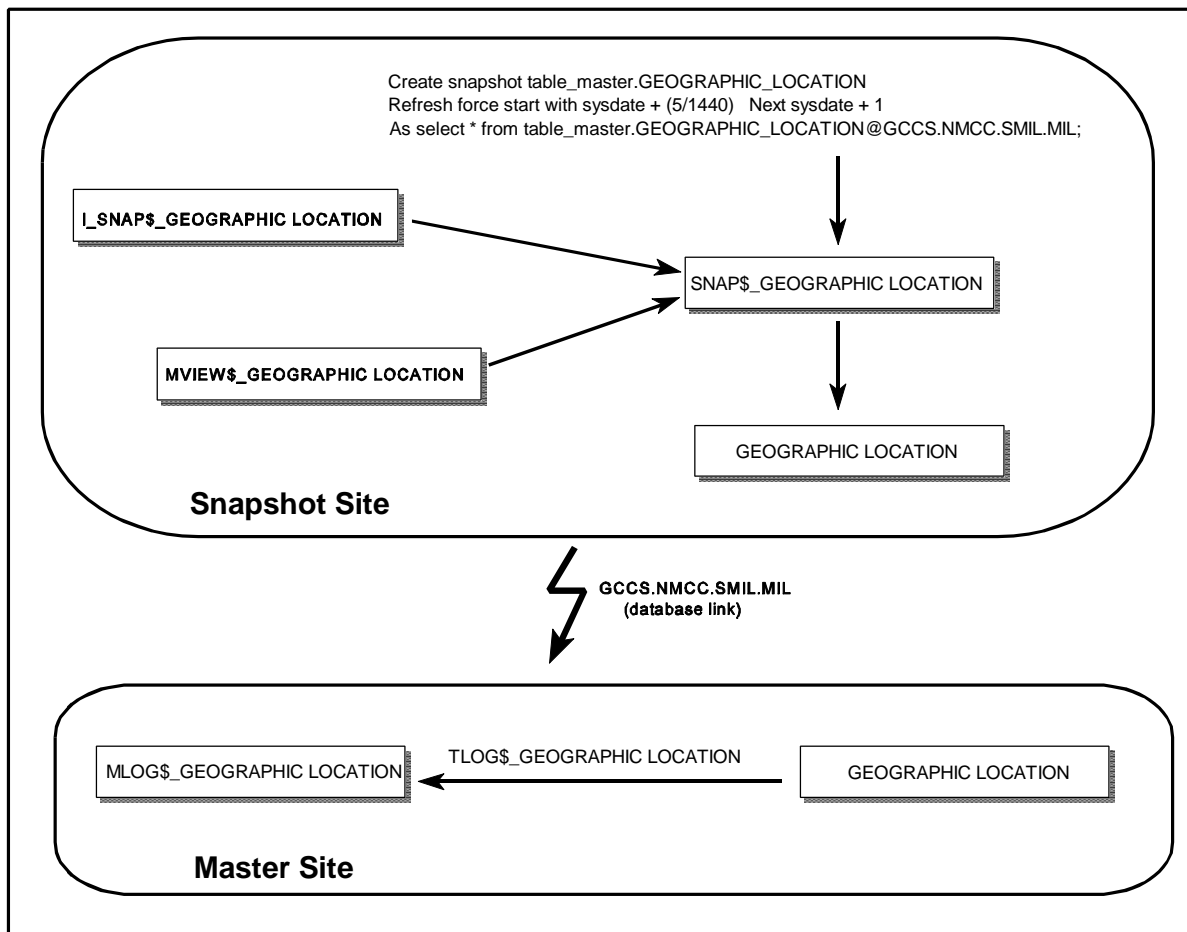


Figure K-1: Example of Snapshot Architecture.

Refer to Subsection K.5 for a listing of snapshot objects found in the GCCS Core Database. Further information on the purpose and creation of each of these objects can be found in the next two subsections.

K.1.2 Snapshot Refresh Process

The snapshot architecture is designed to track all data manipulation performed against the master tables. This information is used by each snapshot site to update their copy of the replicated tables. A snapshot is said to have been refreshed when it has updated the data in its snapshot matching the data in the corresponding master table at NMCC. By default, the snapshots in the GCCS Core Database are set up to be automatically refreshed once per day. Each of the snapshot sites has a unique time of day to do an automated refresh. What hour that is depends on the data found in the table snapshot_site_refresh. See Subsection K.3 for instructions on changing this time setting and executing a manual refresh.

Because of the way the snapshots are defined, Oracle will always do a force refresh. The only time Oracle will do a different type of refresh is if a manual refresh is being performed and something other than a force refresh is specified. During a force refresh, Oracle first tries to do a fast refresh using the snapshot log of the associated master table. The fast refresh updates only rows that have changed since the previous refresh. If Oracle cannot perform the fast refresh, a complete refresh is done. This complete refresh actually re-executes the snapshot's defining query.

During an automated refresh, not all of the snapshot tables are refreshed at once. Instead, three consecutive refreshes are performed. This helps to balance the load on the master site database. One refresh group is done at a time. In order to maintain the relationships between some of the snapshots, all related snapshots are in the same refresh group. This ensures that a snapshot site receives consistent data. The three refresh groups and their associated snapshots are:

C	refresh_group1	AIRPORT
		AIRPORT_AIRCRAFT_CAT
		AIRPORT_CLEARANCE
		AIRPORT_REMARK
		APRON
		GEOGRAPHIC_LOCATION
		HARBOR
		SEAPORT
		SEAPORT_ANCHORAGE
		SEAPORT_BERTH_DEPTH
		SEAPORT_CARGO_CAPACITY
		SEAPORT_CHANNEL
		SEAPORT_CLEARANCE
		SEAPORT_CRAFT
		SEAPORT_CRANE
		SEAPORT_FLOATING_CRANE
		SEAPORT_MHE
		SEAPORT_REMARK
		SEAPORT_STANDARD_BERTH
		SEAPORT_STRUCTURE
		WHARF
		WHARF_BERTH

		WHARF_CONTAINER_STORAGE
		WHARF_EQUIPMENT
C	refresh_group2	PORTS_OF_SUPPORT_L10 SERVICE_RESUPPLY_LOC_L2 SERVICE_RESUPPLY_L3 SERVICE_FUEL_L4 SERVICE_FUEL_RATE_L4 SERVICE_ORGN_APPN_L11 SERVICE_SUPPLY_L4 SERVICE_SUPPLY_RATE_L4 UNIT_TYPE_FUEL_L1 UNIT_TYPE_FUEL_RATE_L1 UNIT_TYPE_SUPPLY_L1 UNIT_TYPE_SUPPLY_RATE_L1
C	refresh_group3	EQUIPMENT_TYPE REFERENCE_FILE_STATUS UNIT_TYPE UNIT_TYPE_CARGO_3RD UNIT_TYPE_CARGO_4TH UNIT_TYPE_REPLACEMENT

K.2 SNAPSHOT CREATION PROCESS

Because different snapshot structures are created on a snapshot site versus the master site, different scripts are run during the creation process. Figure K-2 outlines the scripts and the overall creation process for the master site.

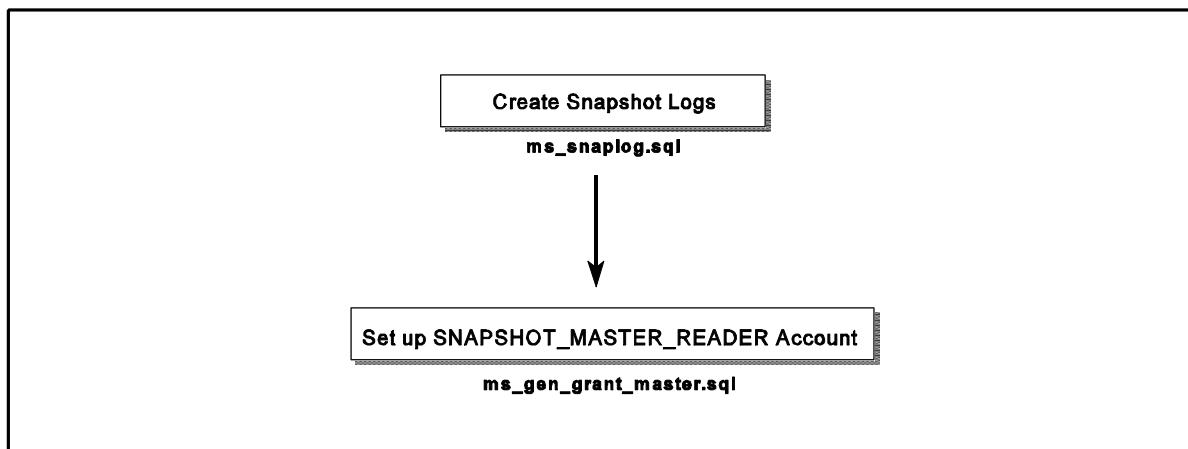


Figure K-2: Creation Process at the Master Site.

The master site needs to run its scripts before the snapshot sites can create their snapshots. At the master site, the first step is to create the snapshot logs. These logs will keep track of any data manipulations that take place on the corresponding master tables. Next, the Oracle account `SNAPSHOT_MASTER_READER` is created and granted select privileges on all the snapshot logs and all the corresponding master tables. For security reasons, this account is granted only minimal permissions to view reference file tables on the NMCC database. This account is later used by the snapshot sites to create their snapshots.

Because the `SNAPSHOT_MASTER_READER` account is always used during a refresh, it is important that NMCC alert all of the other sites of any change to `SNAPSHOT_MASTER_READER`'s password. See Subsection K.3 for further information on what the snapshot sites need to do when this password is changed.

Once the master site is set up, each snapshot site runs the scripts needed to create their snapshots. Figure K-3 outlines the scripts and the overall creation process for a snapshot site.

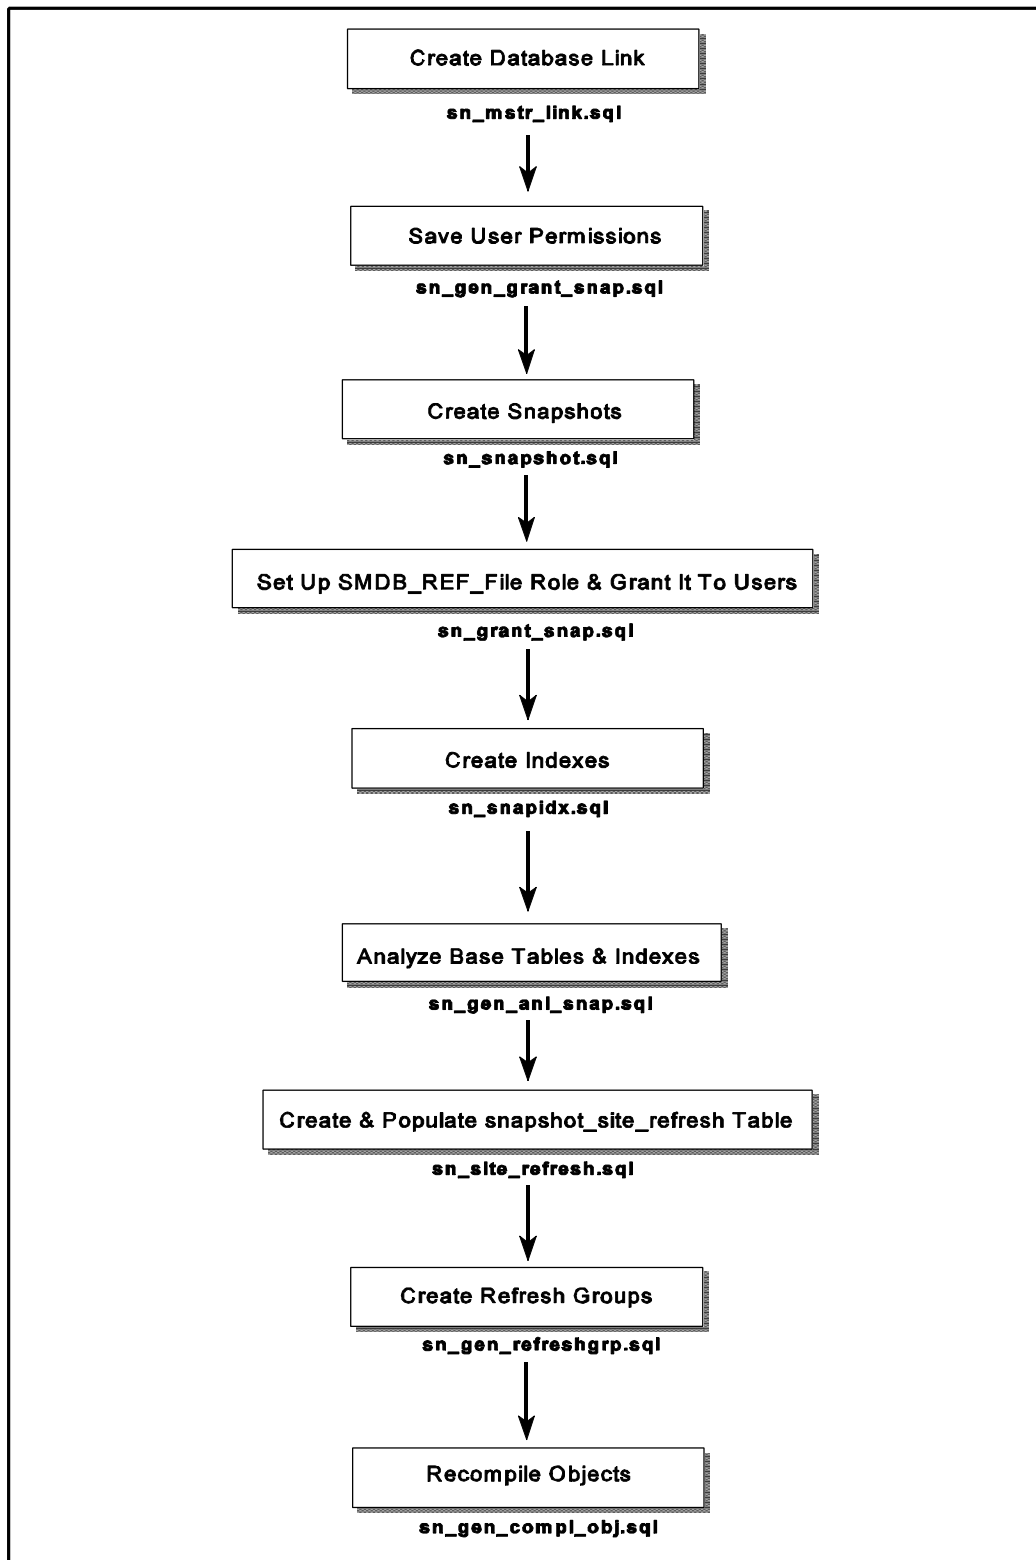


Figure K-3: Creation Process at the Snapshot Site.

K.3 SNAPSHOT ADMINISTRATION

Once the snapshot creation process has been successfully completed, Oracle automatically does most of the day to day snapshot administration. This section details some of the manual tasks that may need to be done. For information on troubleshooting specific problems, refer to the readme text file **SNAP_ADMIN.txt** under the /h/SMDB directory.

K.3.1 chg_rfrsh.sql

A snapshot site can change the hour of the day that its daily automated refresh is scheduled to take place. If the time needs to be changed, use the script **chg_rfrsh.sql** under the /h/SMDB/Scripts/SM_bld_tables directory. When this script is run, the user will be prompted for the new automated snapshot refresh time. This time is based on the local site time for the <oracle> database owner account. The script will then update the automated refresh time for all three refresh groups. The next automated refresh will take place at the new time.

Care should be given when choosing a new automated refresh time. The performance time for each refresh will be at its best when only one snapshot site at a time is communicating with NMCC to do a refresh. In addition, the database load at NMCC may be unacceptable if numerous sites perform refreshes simultaneously.

K.3.2 Reestablishing the Database Link

Every time NMCC changes the password for its SNAPSHOT_MASTER_READER account, all the snapshot sites need to reestablish their database link to NMCC. To do this, each snapshot site must perform the following steps:

- Step 1** Log into Unix as <oradba> at the snapshot site.
- Step 2** Log into sqlplus as <table_master>.
- Step 3** Give the following command at the sqlplus prompt. Enter the new password for SNAPSHOT_MASTER_READER when prompted.

<@/h/SMDB/Scripts/SM_bld_tables/sn_mstr_link>

K.3.3 Manual Refresh

Any snapshot can be manually refreshed at the snapshot site even if it is normally automatically refreshed. There are three types of manual refreshes. The first type refreshes a single snapshot. The second type refreshes all the snapshots associated with a refresh group. The third type refreshes all the snapshots within the database. Refer to the instructions in the subsections below for each type of refresh.

The script **chksnap.sql** in the /h/SMDB/Scripts/SM_bld_tables directory can supply some of the information needed to determine the type of manual refresh that would be appropriate. When this script runs, it reports the date and time when the next scheduled refresh of each refresh group will

take place and when each individual snapshot was last refreshed. Further information can be found within the Oracle data dictionary views **ALL_REFRESH**, **USER_REFRESH**, **DBA_REFRESH**, **ALL_REFRESH_CHILDREN**, **USER_REFRESH_CHILDREN**, and **DBA_REFRESH_CHILDREN**.

K.3.3.1 Manual Refresh of an Individual Snapshot. To refresh an individual snapshot, use the REFRESH stored procedure of the package DBMS_SNAPSHOT. Give the following command as <oradba> or <table_master> at the sqlplus or sqldb prompt where *schema.snapshot_name* is the name of the snapshot to be refreshed and the *refresh_option* is 'F' or 'f' for a fast refresh, 'C' or 'c' for a complete refresh, or '?' for the default force refresh:

```
<exec dbms_snapshot.refresh('schema.snapshot_name','refresh_option');>
```

For example, when <oradba> executes the following SQL statement within sqlplus, a fast refresh of the GEOGRAPHIC_LOCATION snapshot is performed:

```
<exec dbms_snapshot.refresh('table_master.geographic_location','f');>
```

K.3.3.2 Manual Refresh of a Refresh Group. To refresh a snapshot group, use the REFRESH stored procedure of the package DBMS_REFRESH. Give the following command as <oradba> or <table_master> at the sqlplus or sqldb prompt where *schema.refresh_group_name* is the name of the refresh group to be manually refreshed:

```
<exec dbms_refresh.refresh('<schema.refresh_group_name>');>
```

For example, when <oradba> executes the following SQL statement within sqlplus, all the snapshots in the refresh group refresh_group1 are refreshed:

```
<exec dbms_refresh.refresh('table_master.refresh_group1');>
```

This procedure always uses the snapshot's default type of refresh so in this example a fast refresh would be performed. If an error occurs while executing this command, a complete manual refresh of each individual snapshot in the refresh group must be performed.

K.3.3.3 Manual Refresh of All Snapshots. To refresh all the snapshots at a snapshot site, execute a manual refresh for each of the refresh groups. Give the following commands as <oradba> or <table_master> at the sqlplus or sqldb prompt to refresh all the snapshots at the snapshot site:

```
<exec dbms_refresh.refresh('table_master.refresh_group1');>  
<exec dbms_refresh.refresh('table_master.refresh_group2');>  
<exec dbms_refresh.refresh('table_master.refresh_group3');>
```


K.4 SNAPSHOT RELATED REPORTS

This section contains a series of reports describing snapshot related internal objects. These reports are divided into two categories - reports listing those objects that are unique to the master site and reports containing objects that are unique to the snapshot sites. Refer to Appendix J for a description of all the constraints in the GCCS Core Database, including those that are unique to NMCC.